

Exhibit A

Network Working Group
 Request for Comments: 1587
 Category: Standards Track

R. Coltun
 RainbowBridge Communications
 V. Fuller
 Stanford University
 March 1994

The OSPF NSSA Option

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Table Of Contents

1.0 Abstract	1
2.0 Overview	2
2.1 Motivation	2
2.2 Proposed Solution	3
3.0 Implementation Details	5
3.1 The N-bit	5
3.2 Type-7 Address Ranges	5
3.3 Type-7 LSAs	5
3.4 Originating Type-7 LSAs	7
3.5 Calculating Type-7 AS External Routes	8
3.6 Incremental Updates	10
4.0 Originating Type-5 LSAs	10
4.1 Translating Type-7 LSAs	10
4.2 Flushing Translated Type-7 LSAs	13
5.0 Acknowledgements	13
6.0 References	13
7.0 Security Considerations	13
8.0 Authors' Addresses	14
Appendix A: Type-7 LSA Packet Format	15
Appendix B: The Options Field	16
Appendix C: Configuration Parameters	17

1.0 Abstract

This document describes a new optional type of OSPF area, somewhat humorously referred to as a "not-so-stubby" area (or NSSA). NSSAs are similar to the existing OSPF stub area configuration option but have the additional capability of importing AS external routes in a limited fashion.

Coltun & Fuller


[Page 1]

2.0 Overview

2.1 Motivation

Wide-area transit networks (such as the NSFNET regionals) often have connections to moderately-complex "leaf" sites. A leaf site may have multiple IP network numbers assigned to it.

Typically, one of the leaf site's networks is directly connected to a router provided and administered by the transit network while the others are distributed throughout and administered by the site. From the transit network's perspective, all of the network numbers associated with the site make up a single "stub" entity. For example, BARRNet has one site composed of a class-B network, 130.57.0.0, and a class-C network, 192.31.114.0. From BARRNet's perspective, this configuration looks something like this:



where the "cloud" consists of the subnets of 130.57 and network 192.31.114, all of which are learned by RIP on router BR18. Topologically, this cloud looks very much like an OSPF stub area. The advantages of running the cloud as an OSPF stub area are:

1. Type-5 routes (OSPF external link-state advertisements (LSAs)) are not advertised beyond the router labeled "BR10". This is advantageous because the link between BR10 and BR18 may be a low-speed link or the router BR18 may have limited resources.
2. The transit network is abstracted to the "leaf" router BR18 by advertising only a default route across the link between BR10 and BR18.
3. The cloud becomes a single, manageable "leaf" with respect to the transit network.

4. The cloud can become, logically, a part of the transit network's OSPF routing system.
5. Translated type-5 LSAs that are sent into the backbone from the cloud (which is a separate stub area) may be considered "leaf" nodes when performing the Dijkstra calculation.

However, the current definition of the OSPF protocol [1] imposes topological limitations which restrict simple cloud topologies from becoming OSPF stub areas. In particular, it is illegal for a stub area to import routes external to OSPF; it is not possible for routers BR18 and BR10 to both be members of the stub area and to import the routes learned from RIP or other IP routing protocols as type-5 (OSPF external LSAs) into the OSPF system. In order to run OSPF out to BR18, BR18 must be a member of a non-stub area or the OSPF backbone to import routes other than its directly-connected network(s). Since it is not acceptable for BR18 to maintain all of BARRNet's external (type-5) routes, BARRNet is forced by OSPF's topological limitations to run OSPF out to BR10 and to run RIP between BR18 and BR10.

2.2 Proposed Solution

This document describes a new optional type of OSPF area, somewhat humorously referred to as a "not-so-stubby" area (or NSSA) which has the capability of importing external routes in a limited fashion.

The OSPF specification defines two general classes of area configuration. The first allows type-5 LSAs to be flooded throughout the area. In this configuration, type-5 LSAs may be originated by routers internal to the area or flooded into the area by area border routers. These areas, referred to herein as type-5 capable areas (or just plain areas in the OSPF spec), are distinguished by the fact that they can carry transit traffic. The backbone is always a type-5 capable area. The second type of area configuration, called stub, allows no type-5 LSAs to be propagated into/throughout the area and instead depends on default routing to external destinations.

NSSAs are defined in much the same manner as existing stub areas. To support NSSAs, a new option bit (the "N" bit) and a new type of LSA (type-7) area defined. The "N" bit ensures that routers belonging to an NSSA agree on its configuration. Similar to the stub area's use of the "E" bit, both NSSA neighbors must agree on the setting of the "N" bit or the OSPF neighbor adjacency will not form.

Type-7 LSAs provide for carrying external route information within an NSSA. Type-7 AS External LSAs have virtually the same syntax as the

Type-5 AS External LSAs with the obvious exception of the link-state type (see section 3.2 for more details). There are two major semantic differences between type-5 and type-7 LSAs.

- o Type-7 LSAs may be originated by and advertised throughout an NSSA; as with stub areas, NSSA's do not receive or originate type-5 LSAs.
- o Type-7 LSAs are advertised only within a single NSSA; they are not flooded into the backbone area or any other area by border routers, though the information which they contain can be propagated into the backbone area (see section 3.6).

In order to allow limited exchange of external information across an NSSA area border, NSSA border routers will translate selected type-7 LSAs received from the NSSA into type-5 LSAs. These type-5 LSAs will be flooded to all type-5 capable areas. NSSA area border routers may be configured with address ranges so that several type-7 LSAs may be represented by a single type-5 LSA.

In addition, an NSSA area border router can originate a default type-7 LSA (IP address of 0.0.0.0) into the NSSA. Default routes are necessary because NSSAs do not receive full routing information and must have a default route to route to AS-external destinations. Like stub areas, NSSAs may be connected to the backbone at more than one area border router, but may not be used as a transit area. Note that the default route originated by an NSSA area border router is never translated into a type-5 LSA, however, a default route originated by an NSSA internal AS boundary router (one that is not also an area border router) may be translated into a type-5 LSA.

It should also be noted that unlike stub areas, all OSPF summary routes (type-3 LSAs) must be imported into NSSAs. This is to ensure that OSPF internal routes are always chosen over OSPF external (type-7) routes.

In our example topology the subnets of 130.57 and network 192.31.114, will still be learned by RIP on router BR18 but now both BR10 and BR18 can be in an NSSA and all of BARRNets external routes are hidden from BR18; BR10 becomes an NSSA area border router and BR18 becomes an AS boundary router internal to the NSSA. BR18 will import the subnets of 130.57 and network 192.31.114 as type-7 LSAs into the NSSA. BR10 then translates these routes into type-5 LSAs and floods them into BARRNet's backbone.

3.0 Implementation Details

3.1 The N-bit

The N-bit ensures that all members of a NSSA agree on the area's configuration. Together, the N-bit and E-bit reflect an interface's (and consequently the interface's associated area's) external LSA flooding capability. As explained in section 10.5 of the OSPF specification, if type-5 LSAs are not flooded into/throughout the area, the E-bit must be clear in the option field of the received Hello packets. Interfaces associated with an NSSA will not send or receive type-5 LSAs on that interface but may send and receive type-7 LSAs. Therefore, if the N-bit is set in the options field, the E-bit must be cleared.

To support the NSSA option an additional check must be made in the function that handles receiving Hello packet to verify that both the N-bit and the E-bit found in the Hello packet's option field match the value of the options that have been configured in the receiving interface. A mismatch in the options causes processing of the received Hello packet to stop and the packet to be dropped.

3.2 Type-7 Address Ranges

NSSA area border routers may be configured with type-7 address ranges. Each address range is defined as an [address,mask] pair. Many separate type-7 networks may then be represented by in a single address range (as advertised by a type-5 LSA), just as a subnetted network is composed of many separate subnets. Area border routers may then summarize type-7 routes by advertising a single type-5 route for each type-7 address range. The type-5 route, resulting from a type-7 address range match will be distributed to all type-5 capable areas. Section 4.1 gives the details of generating type-5 routes from type-7 address ranges.

A type-7 address range includes the following configurable items.

- o An [address,mask] pair.
- o A status indication of either Advertise or DoNotAdvertise.
- o External route tag.

3.3 Type-7 LSAs: NSSA External Link-State Advertisements

External routes are imported into NSSAs as type-7 LSAs by the NSSA's AS boundary routers. An NSSA AS boundary routers is a router which

has an interface associated with the NSSA and is exchanging routing information with routers belonging to another AS. As with type-5 LSAs a separate type-7 LSA is originated for each destination network. To support NSSA areas, the link-state database must therefore be expanded to contain a type-7 LSA.

Type 7-LSAs are identical to type-5 LSAs except for the following (see section 12.3.4 "AS external links" in the OSPF specification).

1. The type field in the LSA header is 7.
2. Type-7 LSAs are only flooded within the NSSA. The flooding of type-7 LSAs follow the same rules as the flooding of type 1-4 LSAs.
3. Type-7 LSAs are kept within the NSSA's LSDB (are area specific) whereas because type-5 LSAs are flooded to all type-5 capable areas, type-5 LSAs global scope in the router's LSDB.
4. At the area border router, selected type-7 LSAs are translated into type 5-LSAs and flooded into the backbone.
5. Type 7 LSAs have a propagate (P) bit which is used to flag the area border router to translate the type-7 LSA into a type-5 LSA. Examples of how the P-bit is used for loop avoidance are in the following sections.
6. Those type-7 LSAs that are to be translated into type-5 LSAs must have their forwarding address set. Type-5 LSAs that have been translated from type-7 LSAs for the most part must contain a forwarding address. The exception to this is if the translation to a type-5 LSA is the result of an address range match, in which case the type-5 LSA will not contain a forwarding address (see section 4.1 for details). The forwarding address contained in type-5 LSAs will result in more efficient routing to the AS external networks when there are multiple NSSA area border routers. Having the forwarding address in the type-7 LSAs will ease the translation of type-7 into type-5 LSAs as the NSSA area border router will not be required to compute the forwarding address.

If the network between the NSSA AS boundary router and the adjacent AS is advertised into OSPF as an internal OSPF

route, the forwarding address should be the next hop address as is currently done in type-5 LSAs, but unlike type-5 LSAs if the intervening network is not advertised into OSPF as an internal OSPF route, the forwarding address should be any one of the router's active OSPF interface addresses.

Type-5 and type-7 metrics and path types are directly comparable.

3.4 Originating Type-7 LSAs

NSSA AS boundary routers may originate type-7 LSAs. All NSSA area border routers must also be AS boundary routers since they all must have the capability of translating a type-7 LSAs into a type-5 LSAs (see section 3.6 routes for the translation algorithm). NSSA area border routers must set the E-bit (external bit) as well as the B-bit (border bit) in its router (type-1) LSAs (both in the backbone and in the NSSA area).

When an NSSA internal AS boundary router originates a type-7 LSA that it wants to be translated into a type-5 LSA by the NSSA area border router (and subsequently flooded into the backbone), it must set the P-bit in the LS header's option field and add a valid forwarding address in the type-7 LSA.

If a router is attached to another AS and is also an NSSA area border router, it may originate a both a type-5 and a type-7 LSA for the same network. The type-5 LSA will be flooded to the backbone (and all attached type-5 capable areas) and the type-7 will be flooded into the NSSA. If this is the case, the P-bit must be reset in the type-7 NSSA so the type-7 LSA isn't again translated into a type-5 LSA by another NSSA area border router.

A type-7 default route (network 0.0.0.0) may be originated into the NSSA by an NSSA area border router or by an NSSA AS boundary router which is internal to the NSSA. The type-7 default route originated by the NSSA area border router must have the P-bit reset so that the default route originated by the NSSA area border router will not find its way out of the NSSA into the rest of the AS system via another NSSA area border router. The type-7 default route originated by an NSSA AS boundary router which is not an NSSA area border router may have the P-bit set. Type-7 routes which are originated by the NSSA area border router will not get added to other NSSA area border router's routing table.

A default route must not be injected into the NSSA as a summary (type-3) LSA as in the stub area case. The reason for this is that the preferred summary default route would be chosen over all more

specific type-7 routes. Because summary routes are preferred to external routes and to ensure that summary routes are chosen over external within the NSSA, all summary routes (unlike stub areas in which this is optional) must be imported into an NSSA.

3.5 Calculating Type-7 AS External Routes

This section is very similar to section 16.4 (Calculating AS external routes) in the OSPF specification. An NSSA area border router should examine both type-5 LSAs and type-7 LSAs if either type-5 or type-7 routes need to be updated. Type-7 LSAs should be examined after type-5 LSAs. An NSSA internal router should examine type-7 LSAs when type-7 routes need to be recalculated.

In relation to the steps to calculate the routing table as presented in the OSPF specification (chapter 16, "Calculation of the Routing Table"), type-7 LSAs should be examined after step 5 where the routes to external destinations are calculated.

Type-7 routes are calculated by examining type-7 LSAs. Each of LSAs are considered in turn. Most type-7 LSAs describe routes to specific IP destinations. A type-7 LSA can also describe a default route for the NSSA (destination = DefaultDestination). For each type-7 LSA:

1. If the metric specified by the LSA is LSInfinity, the age of the LSA equals MaxAge or the advertising router field is equal to this router's router ID, examine the next advertisement.
2. Call the destination described by the LSA N. Look up the routing table entry for the AS boundary router (ASBR) that originated the LSA. If no entry exists for the ASBR (i.e., ASBR is unreachable), do nothing with this LSA and consider the next in the list.

If the destination is the default route (destination = DefaultDestination) and if the originator of the LSA and the calculating router are both NSSA area border routers do nothing with this LSA and consider the next in the list.

Else, this LSA describes an AS external path to destination N. If the forwarding address (as specified in the forwarding address field of the LSA) is 0.0.0.0, the packets routed to the external destination N will be routed to the originating ASBR. If the forwarding address is not 0.0.0.0, look up the forwarding address in the routing table. Packets routed to the external destination N will be routed within the NSSA to this forwarding address. An intra-area path

must therefore exist to the forwarding address. If no such path exists, do nothing with the LSA and consider the next in the list.

Call the routing table distance to the forwarding address (or the distance to the originating ASBR if the forwarding address is 0.0.0.0) X , and the cost specified in the type-7 LSA Y . X is in terms of the link-state metric, and Y is a Type-1 or Type-2 external metric.

3. Now, look up the routing table entry for the destination N . If no entries exist for N , install the AS external path to N , with the next hop equal to the list of next hops to the forwarding address/ASBR, and the advertising router equal to ASBR. If the external metric type is 1, then the path-type is set to Type-1 external and the cost is equal to $X + Y$. If the external metric type is 2, the path-type is set to Type-2 external, the link-state component of the route's cost is X , and the Type-2 cost is Y .
4. Else, if the paths present in the table are not Type-1 or Type-2 external paths, do nothing (AS external paths have the lowest priority).
5. Otherwise, compare the cost of this new AS external path to the ones present in the table. Note that type-5 and type-7 routes are directly comparable. Type-1 external paths are always shorter than Type-2 external paths. Type-1 external paths are compared by looking at the sum of the distance to the forwarding address/ASBR and the advertised Type-1 paths ($X+Y$). Type-2 external paths are compared by looking at the advertised Type-2 metrics, and then if necessary, the distance to the forwarding address/ASBR.

When a type-5 LSA and a type-7 LSA are found to have the same type and an equal distance, the following priorities apply (listed from highest to lowest) for breaking the tie.

- a. Any type 5 LSA.
- b. A type-7 LSA with the P-bit set and the forwarding address non-zero.
- c. Any other type-7 LSA.

If the new path is shorter, it replaces the present paths in the routing table entry. If the new path is the same cost, it is added to the routing table entry's list of paths.

3.6 Incremental Updates

Incremental updates for type-7 LSAs should be treated the same as incremental updates for type-5 LSAs (see section 16.6 of the OSPF specification). That is, if a new instance of a type-7 LSA is received it is not necessary to recalculate the entire routing table. If there is already an OSPF internal route to the destination represented by the type-7 LSA, no recalculation is necessary. Otherwise, the procedure in the proceeding section will have to be performed but only for the external routes (type-5 and type-7) whose networks describe the same networks as the newly received LSA.

4.0 Originating Type-5 LSAs

4.1 Translating Type-7 LSAs Into Type-5 LSAs

This step is performed as part of the NSSA's Dijkstra calculation after type-5 and type-7 routes have been calculated. If the calculating router is not an area border router this translation algorithm should be skipped. All reachable area border routers in the NSSA should now be examined noting the one with the highest router ID. If this router has the highest router ID, it will be the one translating type-7 LSAs into type-5 LSAs for the NSSA, otherwise the translation algorithm should not be performed.

All type-7 routes that have been added to the routing table should be examined. If the type-7 LSA (associated with the route being examined) has the P-bit set and a non-zero forwarding address, the following steps should be taken.

The translation procedure must first check for a configured type-7 address range. Recall that an type-7 address range consists of an [address,mask] pair and a status indication of either Advertise or DoNotAdvertise. At most a single type-5 LSA is made for each range. If the route being examined falls within the type-7 address range, (the [address,mask] pair of the route equal to or a more specific instance of the [address,mask] pair of the type-7 address range), one of following three actions may take place.

1. When the range's status indicates Advertise and the route's address and mask are equal to the address and mask of the type-7 range a type-5 LSA should be originated if:
 - o there currently is no type-5 LSA originated from this router corresponding to the type-7 LSA,

- o the path type or the metric in the corresponding type-5 LSA is different from the type-7 LSA or
- o The forwarding address in the corresponding type-5 LSA is different from the type-7 LSA.

The newly originated type-5 LSA will describe the same network and have the same network mask, metrics, forwarding address, external route tag and path type as the type-7 LSA, however, the advertising router field will be the router ID of this area border router.

2. When the range's status indicates Advertise and the route's address or mask indicates a more specific route (i.e., the route's address is subsumed by the range or the route has a longer mask), a type-5 LSA is generated with link-state ID equal to the range's address (if necessary, the link-state ID can also have one or more of the range's "host" bits set; see Appendix F of the OSPF specification for details), the network mask, external route tag and path type will be set to the configured type-7 range values. The advertising router field will be the router ID of this area border router. The forwarding address will not be set. The path type should always be set to the highest path type that is subsumed by the net range. The metric for the type-5 LSA will be set as follows:

- o if the path type is external type 2, the type-5 metric should be set to the largest type-7 metric subsumed by this net range + 1.
- o if the path type is external type 1, the type-5 metric should be set to the largest metric.

For example, given a net range of [10.0.0.0, 255.0.0.0] for an area that has type-7 routes of:

```
10.1.0.0 path type 1, metric 10
10.2.0.0 path type 1, metric 11
10.3.0.0 path type 2, metric 5
```

a type-5 LSA will be generated with a path type of 2 and a metric of 6.

As another example, given a net range of [10.0.0.0, 255.0.0.0] for an area that has type-7 routes of:

```
10.1.0.0 path type 1, metric 10
10.2.0.0 path type 1, metric 11
10.3.0.0 path type 1, metric 5
```

a type-5 LSA will be generated with a path type of 1 and a metric of 11.

These metric and path type rules will avoid routing loops in the event that path type 1 and 2 are both used within the area.

3. When the range's status indicates DoNotAdvertise, the type-5 LSA is suppressed and the component networks remain hidden from the rest of the AS.

By default (given that the P-bit is set and the LSA has a non-zero forwarding address) if a network is not contained in any explicitly configured address range, a type-7 to type-5 LSA translation will occur.

A new instance of a type-5 LSA should be originated and flooded to all attached type-5 capable areas if one of the following is true.

1. There currently is no type-5 LSA originated from this router corresponding to the type-7 LSA.
2. The path type or the metric in the corresponding type-5 LSA is different from the type-7 LSA.
3. The forwarding address in the corresponding type-5 LSA is different from the type-7 LSA.

The newly originated type-5 LSAs will describe the same network and have the same network mask, metrics, forwarding address, external route tag and path type as the type-7 LSA. The advertising router field will be the router ID of this area border router.

As with all newly originated type-5 LSAs, a type-5 LSA that is the result of a type-7 to type-5 translation (type-7 range or default case) is flooded to all attached type-5 capable areas.

4.2 Flushing Translated Type-7 LSAs

If an NSSA area border router has translated a type-7 LSA to a type-5 LSA that should no longer be translated, the type-5 LSA should be flushed (set to MaxAge and flooded). The translated type-5 LSA should be flushed whenever the routing table entry that caused the translation changes so that either the routing table entry is unreachable or the entry's associated LSA is not a type-7 with the P-bit set and a non-zero forwarding address.

5.0 Acknowledgements

This document was produced by the OSPF Working Group, chaired by John Moy.

In addition, the comments of the following individuals are also acknowledged:

Phani Jajjarvarpu	cisco
Dino Farinacci	cisco
Jeff Honig	Cornell University
John Moy	Proteon, Inc.
Doug Williams	IBM

6.0 References

- [1] Moy, J., "OSPF Version 2", RFC 1583, Proteon, Inc., March 1994.
- [2] Moy, J., "Multicast Extensions to OSPF", RFC 1584, Proteon, Inc., Proteon, Inc., March 1994.

7.0 Security Considerations

Security issues are not discussed in this memo.

RFC 1587

OSPF NSSA Option

March 1994

8.0 Authors' Addresses

Rob Coltun
RainbowBridge Communications

Phone: (301) 340-9416
EMail: rcoltun@rainbow-bridge.com


Vince Fuller
BARRNet
Stanford University
Pine Hall 115
Stanford, CA, 94305-4122

Phone: (415) 723-6860
EMail: vaf@Valinor.Stanford.EDU

Coltun & Fuller

[Page 14]

Appendix A: Type-7 Packet Format



The definitions of the link-state ID, network mask, metrics and external route tag are the same as the definitions for the type-5 LSAs (see A.4.5 in the OSPF specification) except for:

The Forwarding Address

If the network between the NSSA AS boundary router and the adjacent AS is advertised into OSPF as an internal OSPF route, the forwarding address should be the next hop address but if the intervening network is not advertised into OSPF as an internal OSPF route, the forwarding address should be any one of the router's active OSPF interface addresses.

Appendix B: The Options Field

The OSPF options field is present in OSPF Hello packets, Database Description packets and all link-state advertisements. See appendix A.2 in the OSPF specification for a description of option field.

```

-----
| * | * | * | * | N/P | MC | E | T |
-----

```

The Type-7 LSA options field

- T-bit: The T-bit describes the router's TOS capability.
- E-bit: Type-5 AS external link advertisements are not flooded into/through OSPF stub and NSSA areas. The E-bit ensures that all members of a stub area agree on that area configuration. The E-bit is meaningful only in OSPF Hello packets. When the E-bit is reset in the Hello packet sent out a particular interface, it means that the router will neither send nor receive type-5 AS external link state advertisements on that interface (in other words, the interface connects to a stub area). Two routers will not become neighbors unless they agree on the state of the E-bit.
- MC-bit: The MC-bit describes the multicast capability of the various pieces of the OSPF routing domain [2].
- N-bit: The N-bit describes the the router's NSSA capability. The N-bit is used only in Hello packets and ensures that all members of an NSSA agree on that area's configuration. When the N-bit is reset in the Hello packet sent out a particular interface, it means that the router will neither send nor receive type-7 LSAs on that interface. Two routers will not form an adjacency unless they agree on the state of the N-bit. If the N-bit is set in the options field, the E-bit must be reset.
- P-bit: The P-bit is used only in the type-7 LSA header. It flags the NSSA area border router to translate the type-7 LSA into a type-5 LSA.

Appendix C: Configuration Parameters

Appendix C.2 in the OSPF specification lists the area parameters. The area ID, list of address ranges for type-3 summary routes and authentication type remain unchanged. Section 3.2 of this document lists the configuration parameters for type-7 address ranges.

For NSSAs the external capabilities of the area must be set to accept type-7 external routes. Additionally there must be a way of configuring the NSSA area border router to send a default route into the NSSA using a specific metric (type-1 or type-2 and the actual cost).